Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 10 July 2012 (MN LATEX style file v2.2)

A Sampling Strategy for High-Dimensional Spaces Applied to Free-Form Gravitational Lensing

Mario Lubini⋆ and Jonathan Coles†
Institut für Theoretische Physik, Universität Zürich, Winterthurerstr. 190, 8057 Zürich, Switzerland

arXiv:1207.1722v1 [astro-ph.IM] 6 Jul 2012

ABSTRACT
We present a novel proposal strategy for the Metropolis-Hastings algorithm designed to efficiently sample general convex polytopes in 100 or more dimensions. This improves upon previous sampling strategies used for free-form reconstruction of gravitational lenses, but is general enough to be applied to other fields. We have written a parallel implementation within the lens modeling framework GLASS. Testing shows that we are able to produce uniform uncorrelated random samples which are necessary for exploring the degeneracies inherent in lens reconstruction.

Key words: gravitational lensing: strong - methods: numerical - methods: statistical

1 INTRODUCTION

Some inversion problems in astrophysics make it desirable to search or sample a high dimensional solution domain $S \subset \mathbb{R}^n$ such that S is bounded by the linear constraints

$$\mathbf{A}x \leqslant b \tag{1}$$

where $\mathbf{A} \in \mathbb{R}^{m\times n}$ and b is a constant vector. A classic application is Schwarzschild's construction of triaxial stellar systems in equilibrium (Schwarzschild 1979). Given a three-dimensional discretized target density function ρ_j, the number of stars c_i on a given orbit i is found by solving

$$\rho_j = \sum_i c_i \sigma_{i,j} \tag{2}$$

where $\sigma_{i,j}$ is the orbit density. The orbit density is calculated *a priori* using test particles in a fixed potential corresponding to ρ. However, searching the model space was not feasible at the time and only some particular models were considered. More recent work has further developed this technique (Schwarzschild 1982; Merritt 1999; Cappellari et al. 2006, and references therein).

In this paper we consider applications to gravitational lensing. Lensing has had quite a long history, beginning with the first direct evidence of general relativity, but until 1979 with the discovery of the extra-solar lens Q0957+561 (Young et al. 1981) the field was largely of only theoretical interest (Refsdal 1964b,a). Today more than one hundred strong lensing objects are known with many studied in great detail (e.g., Kochanek et al. 1999; Faure et al. 2008; Auger et al. 2009). Future surveys promise to deliver thousands more.

Of utmost interest is the mass distribution of the lensing object. Characterizing this distribution is important for understanding the properties of galaxies and clusters (Read et al. 2007; Sereno et al. 2010), galaxy formation and evolution (Tortora et al. 2010; Faure et al. 2011), the nature of dark matter (Clowe et al. 2006), as well as estimating cosmological parameters (Bartelmann & Schneider 2001) and the age of the universe (Saha et al. 2006; Oguri 2007; Coles 2008).

Crucially, the equations governing gravitational lensing are linear in the projected mass density κ. As detailed in Section 2, one can discretize κ onto a grid of pixels and solve for physically motivated solutions by imposing constraints in the form of Eq. (1). Several versions of this idea have been developed by Saha & Williams (2004), Coe et al. (2008), and Koopmans (2005).

This free-form approach is more flexible than simple analytic models, which assume a functional form of the mass profile and may unintentionally break degeneracies. However, this creates a large system of linear equations that is highly underconstrained. To understand the range of degeneracies we therefore require a technique that can explore the space of solutions S. One possible technique is to choose a random point x and accept it if x lies in S. This might be a reasonable method in low dimensions n, as is done in Monte Carlo integration, but the probability of acceptance rapidly approaches zero as n increases. Each of the pixels in the discretization of κ represents one dimension and typically n is greater than 100. Complex systems, where multiple lenses are used, can easily have more than 1000 dimensions. The priors can also be arbitrary, so the simplex will have a very complex shape, although by construction it will always be convex.

General sampling of probability distributions has been

⋆ E-mail: lubini@physik.uzh.ch
† E-mail: jonathan@physik.uzh.ch

© 0000 RAS

a topic of statistics research for many years (e.g., Chib & Greenberg 1995; Robert & Casella 2005). In the case of lensing, the PixeLens algorithm (Saha & Williams 2004) is frequently used. We show, however, that the sampling of this algorithm is not uncorrelated. We address these details and related issues in Section 3 and suggest an alternative based on the Metropolis-Hastings algorithm in Section 4. In Section 5 we discuss the implementation and demonstrate in Section 6 that even for high dimensions we are able to sample our solution space to achieve a uniform uncorrelated random sample. We also achieve significant speed improvements over PixeLens. In Section 7 we discuss future work and applications.

2 FRAMEWORK

There are two primary equations in gravitational lensing (Blandford & Narayan 1986; Schneider et al. 1992; Schneider 2006). The lens equation

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \nabla\psi(\boldsymbol{\theta}) \tag{3}$$

maps an observed position $\boldsymbol{\theta}$ to an unobservable source position $\boldsymbol{\beta}$ through the potential

$$\psi(\boldsymbol{\theta}) = \frac{1}{\pi}\int_{\mathbb{R}^2} \kappa(\boldsymbol{\theta}')\ln|\boldsymbol{\theta}-\boldsymbol{\theta}'|\mathrm{d}\boldsymbol{\theta}' \tag{4}$$

where κ is the dimensionless projected mass density of the lensing object. The Fermat potential

$$\tau(\boldsymbol{\theta}) = \frac{1}{2}|\boldsymbol{\theta}-\boldsymbol{\beta}|^2 - \psi(\boldsymbol{\theta}) \tag{5}$$

measures, up to an affine transformation, the time a photon takes to travel from the source to the observer. Eq. (3) corresponds to the stationary points of Eq. (5). If the source varies in brightness and the arrival times are different for different images then one can measure the physical time delay Δt_{21} between the light curves of $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$.

As in Saha & Williams (2004), we discretize κ into grid cells, or pixels, centered on the lensing object and construct a system of linear equations from $\Delta\tau \propto \Delta t$ and Eq. (3):

$$\mathbf{C}x = d \tag{6}$$

where $\mathbf{C} \in \mathbb{R}^{p\times k}$, d is a constant vector, and x is a vector containing the free parameters κ and $\boldsymbol{\beta}$. These equations only serve to reduce the dimension of the problem k by the number of equalities p since in general $p \leqslant k$, where p is equal to twice the number of observed images plus the number of measured time delays. This reduction is performed with the orthogonal projection

$$\mathrm{proj}(x) = (\mathbf{1} - \mathbf{C}^+\mathbf{C})x + \mathbf{C}^+ d \tag{7}$$

which takes a point x to the solution set of Eq. (6). The matrix $\mathbf{C}^+$ is the Moore-Penrose pseudoinverse of $\mathbf{C}$. A basis of this affine space is given by those eigenvectors of $(\mathbf{1} - \mathbf{C}^+\mathbf{C})$ with eigenvalue equal one. We can therefore, without loss of generality, take this reduced space of dimension $n = k - p$ to be our problem domain.

Since the space is unbounded, we must impose constraints in the form of Eq. (1) to limit ourselves to reasonable and physical solutions. These constraints may derive from data, such as arrival time order or image parity, or from Bayesian priors. We consider only modest priors, such as the mass must be positive everywhere, variations in κ must be smooth, and the local density gradient must point within 45° of the center. A complete discussion can be found in Coles (2008). Our choice of constraints constructs a non-empty compact solution space S, which is a convex polytope, or simplex. By our definition of the solution space, the Bayesian posterior distribution is

$$P(x) \propto \begin{cases} 1 & \text{if } x \in S \\ 0 & \text{if } x \notin S \end{cases} \tag{8}$$

since all $x \in S$ are equally probable. We are interested in an uncorrelated uniform random sample drawn from S, which we will simply refer to as a random sample.

3 REVISITING THE PIXELENS ALGORITHM

Earlier work used the program PixeLens (Saha & Williams 2004) to model gravitational lenses and estimate the Hubble Time H_0^{-1}. The sampling strategy employed in PixeLens is a type of random walk explained in detail in Saha & Williams (2004) and Coles (2008). Here we summarize the algorithm and discuss some problems.

To build a set of sample points $X = \{x_1, x_2, \dots\}$ one begins by selecting a set of vertices $V = \{v_0, v_1, \dots\}$ of S. The first sample point x_1 is chosen uniformly from the chord connecting v_0, v_1. Each new point x_i with $i > 1$ is chosen randomly and uniformly from the chord from v_i through x_{i-1} to the boundary of S. In the limit of infinite samples this algorithm will explore the entire simplex.

To construct V, PixeLens uses the simplex algorithm (Dantzig 1963; Press et al. 2007). This algorithm was designed to maximize (or equivalently minimize) a linear objective function $f(x)$ subject to linear constraints as in Eq. (1) by moving from vertex to vertex of the simplex in a direction that always increases f. It is a standard algorithm in the field of linear programming, where the vertex that maximizes f is the desired result. Finding a particular vertex is not the goal of PixeLens and so each vertex v_i is found by maximizing a random objective function $f(x) = c \cdot x$ where $c = (c_1, \dots, c_n)$ is a random vector with uniform $c_i \in [-1, 1]$.

One issue is that randomly choosing an objective function does not randomly choose a vertex. If the simplex is not a regular polytope there will be some vertices that are chosen with a higher probability than others. This is demonstrated in Figure 1 with a simplex in seven dimensions with 32 vertices. The vertices have been enumerated and sorted by the number of times they were chosen. Clearly some vertices are highly preferred. Even in high dimensions where the choice of a particular vertex is unlikely to occur again, vertices that are particularly acute will be more likely. By not choosing vertices at random the algorithm prefers some regions over others which leads to correlations in the final sample and not all directions will be adequately explored. Vertex selection is also not invariant under some general coordinate transformation $x \to \mathbf{T}x$ for an invertible matrix $\mathbf{T} \in \mathbb{R}^{n\times n}$. Even a change in units may affect the sample distribution and therefore the inferred physical parameters.

Correctly choosing vertices at random is itself a difficult problem. There exist several methods to enumerate the vertices (e.g., Avis & Fukuda 1992; Dyer 1983) but unfortunately the number of vertices has a huge combinatorial

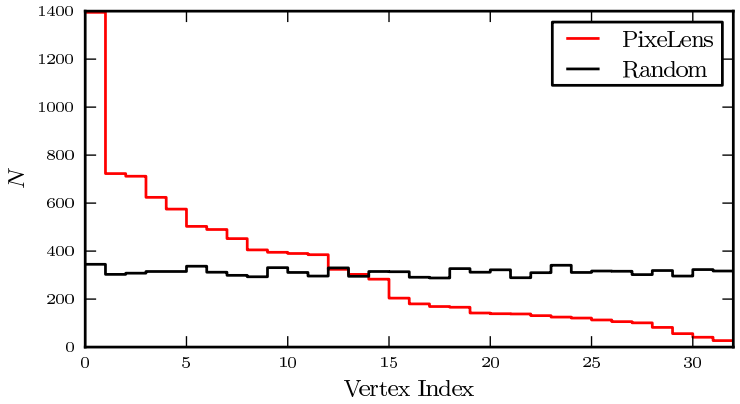

Figure 1. In the PixeLens sampling some vertices are preferred to others when choosing a random objective function. This is shown here using a simplex in seven dimensions with 32 vertices. The red curve is the number of times N a particular vertex was chosen using PixeLens. For comparison, a purely random choice of the same vertices produces the black curve.

upper bound of

$$\binom{m+\lfloor \frac{1}{2}n \rfloor}{m} + \binom{m+\lceil \frac{1}{2}n \rceil - 1}{m} \tag{9}$$

where m is the number of inequalities (McMullen & Shephard 1971).

To avoid the enumeration we modified the simplex algorithm to randomly walk between the vertices. This dispenses with the objective function and simply selects a neighboring vertex to which to move. While this does improve the sampling of the vertices (see Figure 1), it is not without its own problems. If a vertex has many close neighbors, as is likely in high dimensions, the random walk will tend to stay in one region before moving large distances. To compensate, one must run for a long time. The process of moving to a new vertex is computationally costly, however, and incurs numerical error that quickly dominates after too many iterations.

Another issue is that the PixeLens algorithm is based on the vertices of S. The algorithm produces samples that do not follow the target probability distribution function $P(x)$, even if the vertices are chosen randomly. In Figure 2 we demonstrate this for a 100 dimensional hypercube and n-ball, where the vertices have been randomly selected *a priori* to avoid the PixeLens vertex selection algorithm. In the case of the n-ball we have chosen a random set of points on the surface to be the vertices. Points chosen from the hypercube tend to lie in the corners, while points from the n-ball are more closely clustered in the center. In general, the points tend to clump along the chords connecting vertices. In both cases the PixeLens sampling is markedly different than a random sample, although the means are nearly identical.

The sample should be uniformly distributed in S in order to be able to perform statistical analysis on it. For this reason, we chose to explore an alternative method based on a random walk that does not depend on the vertices of the simplex.

4 A NEW MCMC PROPOSAL DENSITY

The Metropolis-Hastings algorithm (Metropolis et al. 1953; Hastings 1970) is a well known method to sample the probability distribution $P(x)$ by generating a Markov chain $X = \{x_1, x_2, \dots\}$. A sample x' is selected from a proposal density function $Q(x', x_i)$ given the current sample x_i and if

$$\alpha < \frac{P(x')Q(x_i, x')}{P(x_i)Q(x', x_i)} \tag{10}$$

where $0 \leqslant \alpha \leqslant 1$ is chosen from a uniform distribution, then x' is accepted and $x_{i+1} = x'$. If x' is rejected the current point is duplicated as $x_{i+1} = x_i$. We will assume that Q is symmetric, i.e., $Q(x', x_i) = Q(x_i, x')$, so that the chance of moving from x_i to x_{i+1} is the same as moving from x_{i+1} to x_i.

One possibility for Q is to simply move by an arbitrary amount in a random direction but the chain may become trapped in narrow regions, especially in high dimensions. To account for the shape of S, Q is often taken to be a multivariate Gaussian distribution $\mathcal{N}(x_i, \widehat{\mathbf{\Sigma}})$, where $\widehat{\mathbf{\Sigma}}$ is the covariance matrix of the sample X. This matrix is an estimate of the covariance matrix $\mathbf{\Sigma}$ of S and can be progressively calculated as the chain is built. Such adaptive chains are no longer Markovian because the reversibility is broken, but one can run an initial adaptive burn-in phase before beginning the Markovian chain with $\widehat{\mathbf{\Sigma}}$ fixed at its last value.

Selecting x' from $\mathcal{N}(x_i, \widehat{\mathbf{\Sigma}})$ is equivalent to selecting y' from the distribution $\mathcal{N}(0, \mathbf{1})$[1] and setting $x' = x_i + \mathbf{E}\,\mathbf{\Lambda}^{1/2} y'$, where $\mathbf{E}$ is the matrix of the eigenvectors e_j of $\widehat{\mathbf{\Sigma}}$ and $\mathbf{\Lambda}$ is the diagonal matrix with the corresponding eigenvalues λ_j. In other words, we move along a randomly selected direction accounting for the shape of S through the eigenvalues.

For a reasonable estimate of $\mathbf{\Sigma}$, particularly in high dimensional spaces, it is important to have $\gg n$ points. When only $\gtrsim n$ points are known, some of the eigenvalues of $\mathbf{\Sigma}$ are strongly underestimated simply due to poor sampling of the

[1] The probability density function for y' is $f(y') = (2\pi)^{-\frac{n}{2}} e^{-\frac{1}{2}\|y'\|^2}$.

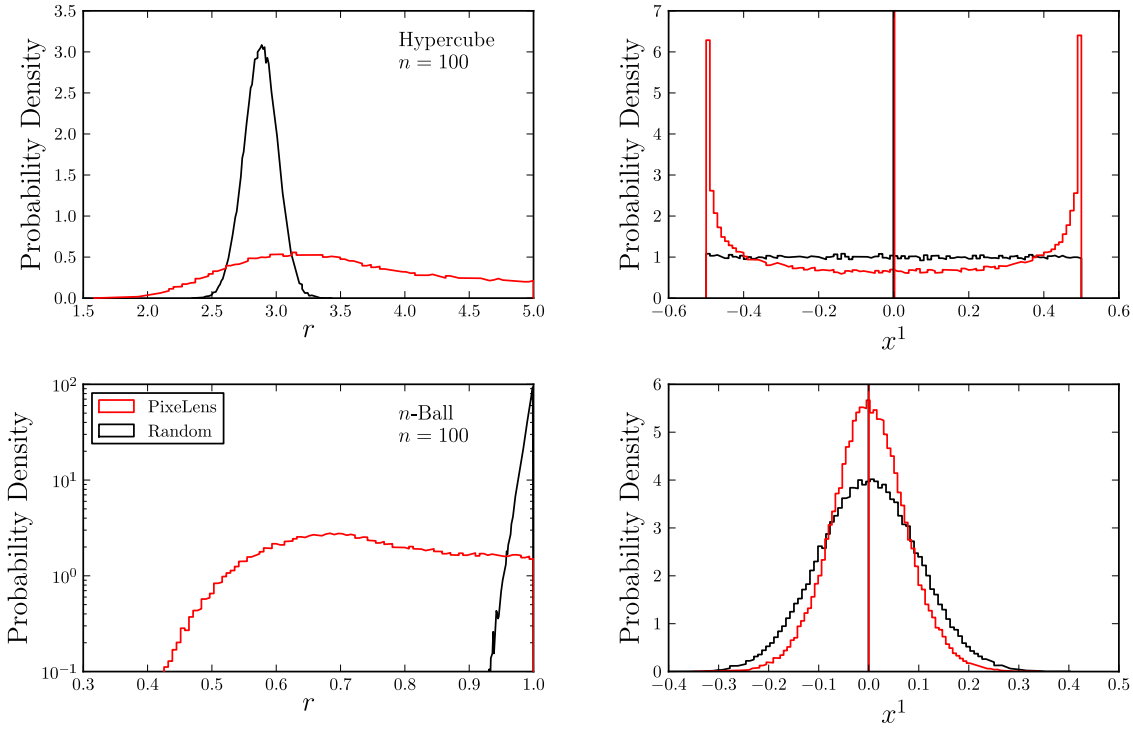

Figure 2. A comparison of sample sets X_P and X_R for the PixeLens sampling algorithm **(red)** and random sampling **(black)**, respectively. Each set has 100,000 items. The means of the plotted distributions (vertical lines) for both the hypercube and n-ball samples are nearly identical, whereas the deviations are not. **(Top left)** We take the sample space S to be a unit cube centered at the origin in 100 dimensions and show the probability distribution function of $r = |x|$ for all $x \in X$. **(Top Right)** The PDF for the single coordinate x^1 of each sample point. Marginalizing over the other coordinates we expect the probability to be uniformly one. Points from X_P clump along the chords connecting vertices. **(Bottom left)** Similarly, we plot the PDF where S is a n-ball. **(Bottom right)** Marginalizing out a single coordinate we see that a random sample has a broader distribution.

space. Even if true random samples were to be drawn directly from S, the shape would be incorrectly estimated. In Figure 3 all the eigenvalues of a 100 dimensional cube should be equal, but for small sample sizes $|X|$ this is clearly not the case. Poorly estimated eigenvalues cause the standard proposal density to undersample S in the direction of the corresponding eigenvectors, even as new points are added to the chain; the new points reinforce the bias that was present in the original $\widehat{\mathbf{\Sigma}}$.

The key improvement from this paper is to use the constraint information of Eq. (1) to hint at the shape of S and achieve a reasonable proposal density despite having a small sample size. We do this in the following way. Let $X = \{x_1, x_2, \ldots, x_k\}$ with $k \geqslant n+1$ be a set of points in $S \subset \mathbb{R}^n$. These points may be chosen in any fashion, but the sample covariance matrix $\widehat{\mathbf{\Sigma}}$ must be invertible, i.e., all x_i's do not lie on the same hyperplane of $\mathbb{R}^n$, and thus the set of the eigenvectors of $\widehat{\mathbf{\Sigma}}$ is an orthonormal basis of $\mathbb{R}^n$. Since S is convex, the mean $\langle X \rangle$ will also be in S. Extending the eigenvector e_j from $\langle X \rangle$ intersects the boundary of S at two points: one in the positive and one in the negative direction. The distance d_j between each pair of boundary points is taken as an estimate of the size of S along e_j.

Our modification takes Q to be the multivariate Gaussian distribution $\mathcal{N}(x_i, \widetilde{\mathbf{\Sigma}})$, where $\widetilde{\mathbf{\Sigma}} = \mathbf{EDE}^T$ and $\mathbf{D}$ is the diagonal matrix of the new $\sigma_j^2 = d_j^2/12$. We therefore select $x' = x_i + \mathbf{E}\,\mathbf{D}^{1/2} y'$. The ellipsoidal shape of Q is thus adjusted by substituting $\mathbf{\Lambda}$ with $\mathbf{D}$ to better approximate the shape of S. In Figure 4 we show schematically the modification. The initial set of points inadequately samples the horizontal direction and therefore the eigenvalues λ_1, λ_2 of the covariance matrix are small in that direction. The new distances d_1, d_2 are a much better approximation and are taken in place of the eigenvalues. This strategy encourages the movement along directions that have been poorly sampled and therefore have a very small variance. This is an important distinction to so-called Hit-and-Run algorithms (Belisle et al. 1993) that step in a random direction and may require many iterations to move through narrow spaces.

Metropolis algorithms in high dimensions can often be inefficient. If the average step length is too big, almost all proposals will fall into low probability regions and be rejected, whereas if the step length is too small, almost all proposals will be accepted but sampling the space will be very slow. The optimal is somewhere in between and can be reached if we regulate the step length by multiplying the σ_i's with some scaling constant c such that the acceptance rate is roughly 25% (Gelman et al. 1996). In our case $P(x)$ is not a multivariate Gaussian distribution and its shape varies from case to case and thus there is no single c. However, we found that for $c \sim 1/n$ the acceptance rate remains reasonable.

Assuming $c = 4/n$, the random walk will typically translate to an average distance of $t_1 = 4\left(\sum_j \sigma_j^2/n^2\right)^{1/2}$ after one step and $t_s = t_1\sqrt{s}/2$ after s steps assuming an

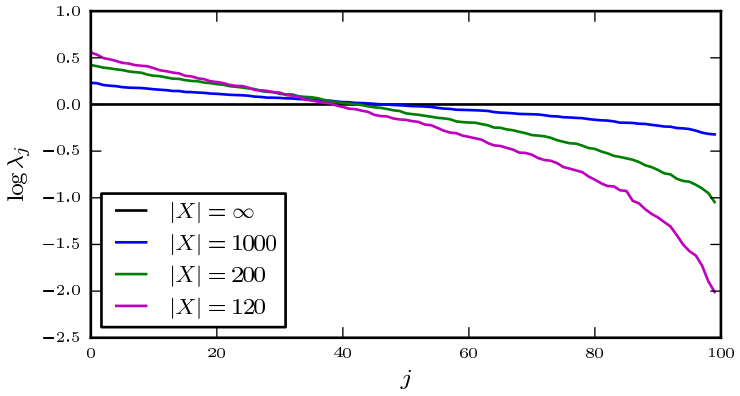

Figure 3. The sorted eigenvalues of the sample covariance matrix of three sets of random points in a 100 dimensional cube with side length $\sqrt{12}$. The set size is $|X|$. With infinite samples we expect all eigenvalues to be equal to unity, but due to the high dimensionality, a random sample does not fully estimate each direction, especially for low $|X|$.

acceptance rate of 25%. In order to produce two uncorrelated points in S, we must make $s = N_s$ steps such that the average traveled distance is of the order of the simplex diameter D. In the case of the hyperrectangle or similar shaped simplices

$$D \sim 2\sqrt{\sum_j \sigma_j^2} = \frac{n}{2} t_1 = t_{n^2} \tag{11}$$

and therefore $N_s = O(n^2)$ steps are needed. In other cases, such as a regular n-simplex, where the approximation for D in Eq. (11) does not hold, $N_s > O(n^2)$ steps may be required. This is not a typical scenario however in lensing, as we demonstrate in Section 6.

The strength of this algorithm is that it is not sensitive to the dispersion of the starting set of points X but only to its mean $\langle X \rangle$. When $\langle X \rangle$ is not a good estimate of the center of S, the algorithm will need some time to remove the starting bias.

5 IMPLEMENTATION

We have implemented our modified sample strategy in a new gravitational lens modeling framework called GLASS. This framework is specifically designed for free-form lens modeling and to allow for easy modification of modeling strategies and priors. Furthermore, we are able to immediately test the implementation by comparing with lensing theory and published results.

As discussed in the previous section, the proposal density Q depends on an estimate for the size of the simplex S. We estimate the size by measuring the distances d_j from the current sample chain mean $\langle X \rangle$ to the boundary of S following the estimated eigenvectors e_j. These diameters are best estimated if $\langle X \rangle$ is close to the simplex mean $\langle S \rangle$ and the vectors e_j are aligned with the true eigenvectors of S.

As is often done (Press et al. 2007), the Markov chain is restricted to move coordinate-wise along the eigenvectors e_j, rather than in a random direction. We rotate S such that these eigenvectors coincide with the standard basis. This provides a significant performance improvement since only one coordinate needs to be updated. In addition, the constraints Eq. (1), typically numbering a few times n in lensing, must only be checked in one coordinate. After N_s steps the last point is rotated back into the original coordinate system and appended to X. With $N_s = O(n^2)$ the running time to produce one sample is $O(n^3)$. While PixeLens also has a theoretical running time of $O(n^3)$ our implementation has a reduced scaling constant resulting in significant performance gains.

Our implementation begins by finding the point x_0 where the temporary variable t is maximized subject to $\mathbf{A}x_0 + t \leqslant b$. For this we use the simplex algorithm but any linear programming algorithm will suffice. The point x_0 is inside S and in some sense "far" from the boundaries.

Initially, the chain walks along the eigenvectors of the matrix $\mathbf{1} - \mathbf{C}^+\mathbf{C}$. The chain is run for a burn-in phase where we collect N_b samples. We typically let $N_b = 10n$. After the first $2n$ samples, and subsequently after each $N_b/10$ samples, we updated Q by calculating $\widetilde{\mathbf{\Sigma}}$ and then continue walking along the respective eigenvectors. The scaling constant c is adjusted to hold the acceptance rate around 25%.

After the burn-in phase, we fix $\widetilde{\mathbf{\Sigma}}$ and c at their final values and run a new chain for as many samples as are desired. In both phases we can run several Markov chains in parallel as long as we ensure that all threads use the same $\widetilde{\mathbf{\Sigma}}$. We have tested this on a shared memory machine using up to 48 CPUs.

As we move to higher dimensional spaces we expect that accumulation of round-off error will tend to produce departures from the equality constraints in Eq. (6). To remove this numerical error we ensure that a sample point x lies on the simplex by using Eq. (7) to project it back onto the simplex.

6 SAMPLING EVALUATION

The stationary distribution of the sample set X from any general MCMC strategy will be the target probability distribution $P(x)$. This is only reached though for a sample size much greater than n. In practice this is not feasible in high dimensions and we must limit our sample size to

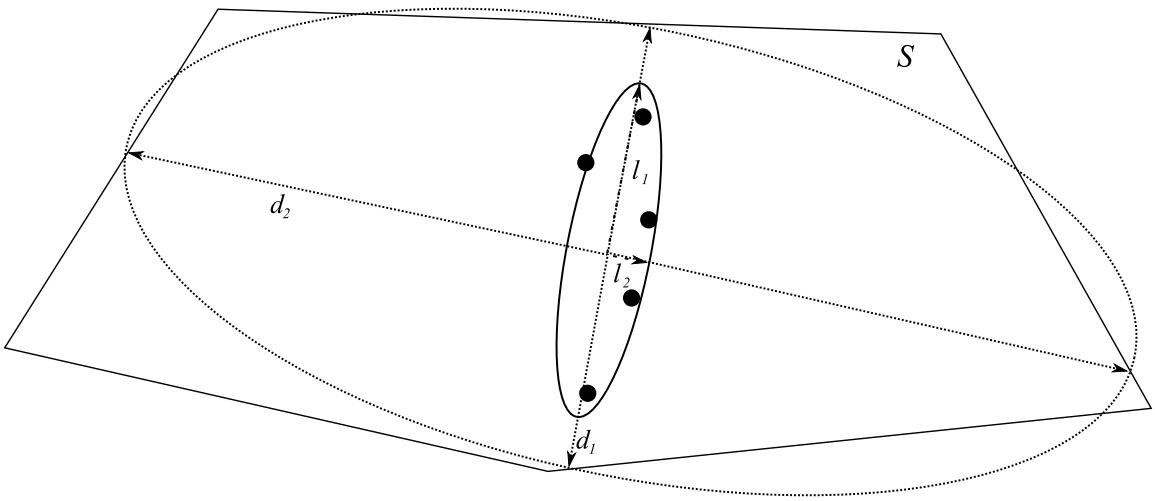

Figure 4. Our modification to the choice of proposal density Q is show schematically. The eigenvectors and corresponding eigenvalues $\lambda_1 = l_1^2, \lambda_2 = l_2^2$ of the covariance matrix of an initial set of points do not accurately capture the shape of the simple S. This is a problem only for high dimensions. We show a two dimensional drawing for simplicity. Since the boundaries of S are known, we extend the eigenvectors until we reach the boundary. We substitute the eigenvalues with $d_1^2/12, d_2^2/12$.

$\gtrsim n$. As we demonstrated in Figure 3 the eigenvectors of a small random sample will not be able to fully describe the solution space but for lensing statistics this is sufficient as we typically marginalize over many parameters. As we want to obtain a uniform uncorrelated random sample, we compare our MCMC implementation to a random sample from a hyperrectangle and a regular n-simplex in 100 dimensions, while varying N_s to measure convergence.

The hyperrectangle

$$H_n = \{(x_1, x_2, \ldots, x_n) \in \mathbb{R}^n \mid 0 \leqslant x_j \leqslant 1/j\} \qquad (12)$$

is straightforward to sample directly and the n-simplex

$$S_n = \{(x_1, x_2, \ldots, x_{n+1}) \in \mathbb{R}^{n+1} \mid 0 \leqslant x_j, \sum_j x_j = 1\} \quad (13)$$

is only slightly more involved[2]. We expect the simplex of a real lens system to be similar to H_n since it has $2n$ inequalities and 2^n vertices. In addition, we also test S_n as an extreme example. We repeated the tests of PixeLens from Section 3 using the hypercube and S_n with $N_s = n^2$ and $N_s = n^{2.5}$, respectively, and show the results in Figure 5. We are able to match the expected distributions of a random sample perfectly.

We test the global properties of our samples X_H, X_S against the random samples R_H, R_S by comparing the eigenvalues of the respective sample covariance matrices. Each sample set contains 1000 points. In Figure 6 we show the sorted eigenvalues for these samples. As expected from Eq. (11) the hyperrectangle converges at $N_s \sim n^2$. The n-simplex requires a larger N_s, as mentioned in Section 4, because the estimate of the diameter from Eq. (11) is no longer valid. In the right panels, we have taken several hundred random sample sets and plotted the 1σ deviations. The plots are normalized to the mean of these random sets. The volume of a simplex can be well approximated by $(\det \widehat{\mathbf{\Sigma}})^{1/2} = \prod_j \sqrt{\lambda_j}$. Table 1 shows the volumes of all the computed samples. The uncertainties have been calculated from 1000 independent runs. Our sampling strategy is in excellent agreement with random samples.

While the eigenvalues paint a global picture, we also tested the local properties of our sample. In particular, we looked at the distribution of nearest neighbor distances. In Figure 7 we show this distribution for H_n and S_n and for increasing N_s. A misalignment with a random sample or multiple peaks are indicators of a correlated sample, such as clumped points. The chains with too low N_s are unable to traverse across the simplex resulting in sample points which are too close to each other. By the time the eigenvalues converge the local distribution also converges.

Finally, we tested our implementation with a simulated triaxial lens mass. The four image positions and respective time delays were calculated using a root finding algorithm built into GLASS. We supplied the value of H_0, all the time delays, and all the image positions to the algorithm without error bars to test the effectiveness of the method. The problem, however, still remains heavily underconstrained. In the near future, we will explore in detail the effects of relaxing these assumptions in a variety of different lens systems.

For the reconstruction, we used a grid of 225 pixels but we assumed radial symmetry to reduce the number of independent pixels to 113. Together with the unknown source position the problem lies on a 104 dimensional affine space. The reconstructed average arrival time surface and images are shown in the left panel of Figure 8. The right panel compares the inferred surface density profile with uncertainties (black error bars) to the original lens profile. The constraint information on the mass profile is contained within the image positions (vertical lines) and therefore the pixels outside are not expected to be well fit. In general, though, this reconstruction is excellent where the information content is highest. As discussed in Section 1, by sampling the solution space, we are able to explore the degeneracies which sim-

[2] To generate a uniform random point in the regular n-simplex S_n choose an $(n+1)$-vector V of i.i.d. numbers drawn from an exponential distribution. Then $V' = V/\sum_j V_j$ is a point in S_n (Devroye 1986).

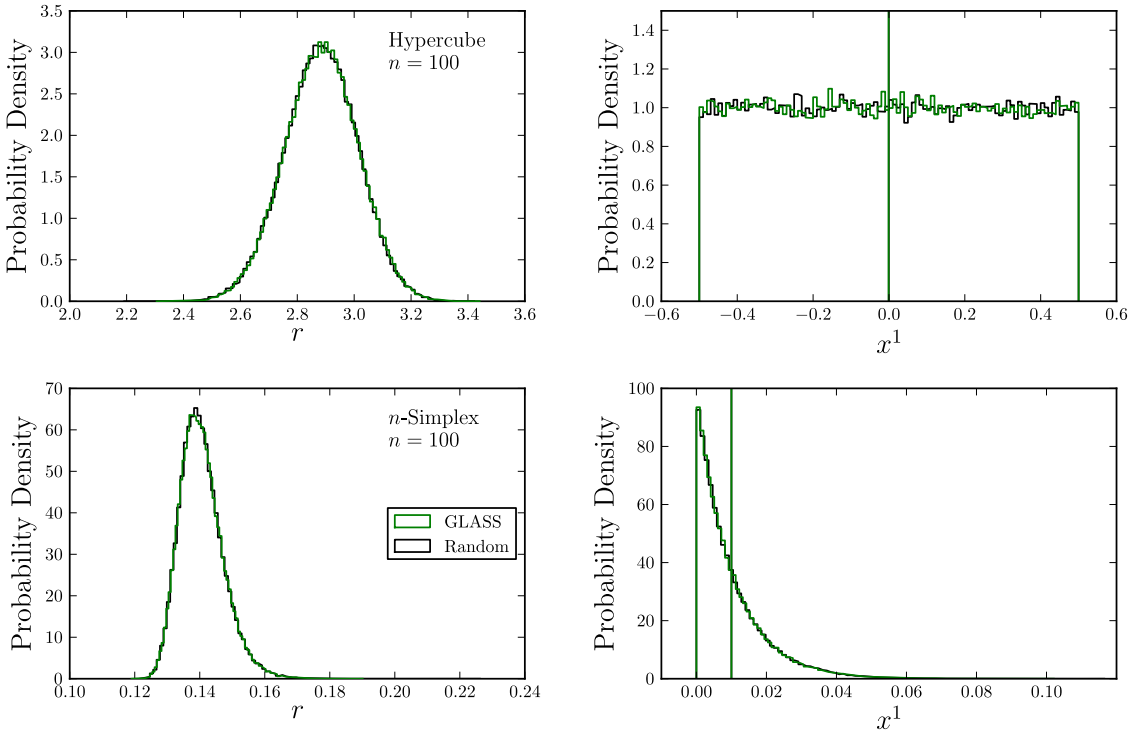

Figure 5. Repeating the tests shown in Figure 2 using GLASS. With the MCMC implementation we correctly recover the same distributions as a random sample. **(Left plots)** The distance $r = |x|$ of each point from the origin. **(Right plots)** The distribution of the first coordinate of $x \in X$, marginalizing over the others. The n-Simplex lies on only one side of the origin. The vertical lines mark the mean of each sample.

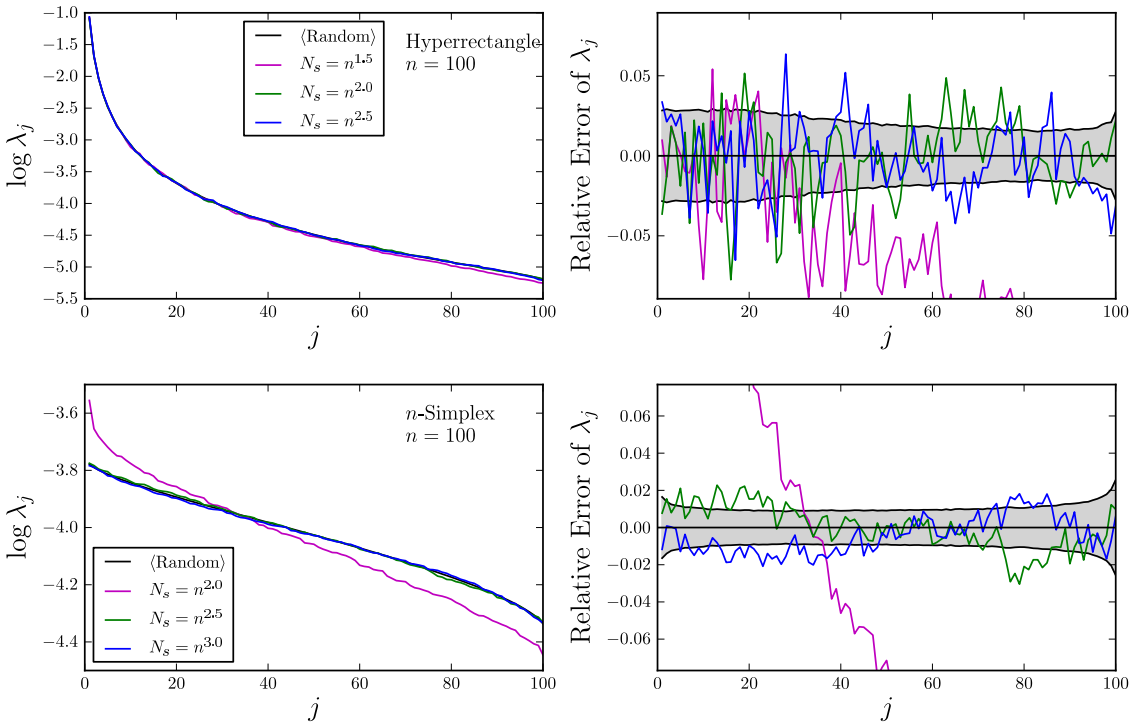

Figure 6. A convergence test of the MCMC sample eigenvalues λ_j with increasing N_s. For comparison, several hundred random sets were generated. The average of the random sets is shown in the left plots. The right plots normalize the MCMC eigenvalues to the mean of the random values. The grey band shows the 1σ deviations of the random sets. The hyperrectangle **(top)** converges by $N_s \sim n^2$ as predicted from Eq. (11), but the n-Simplex **(bottom)** with its different shape requires a larger $N_s \sim n^{2.5}$.

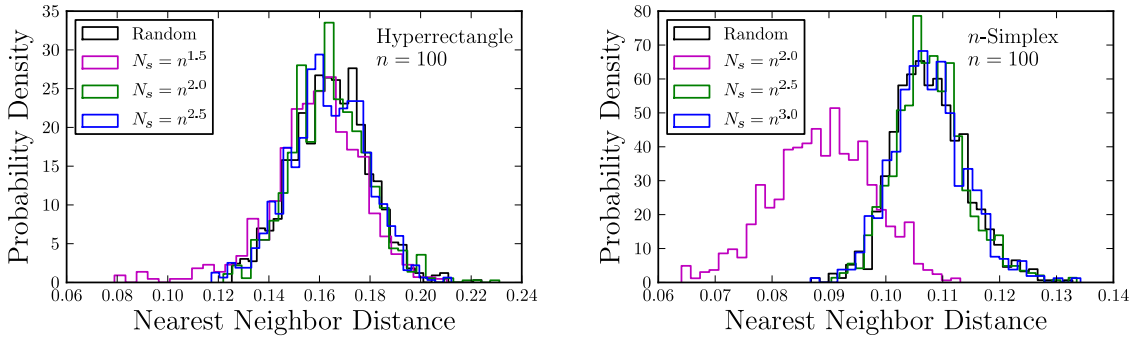

Figure 7. The probability that the nearest neighbor of a sample point is at a given distance in the hyperrectangle **(left)** and n-simplex **(right)**. The histograms produced by GLASS show similar convergence as in Figure 6.

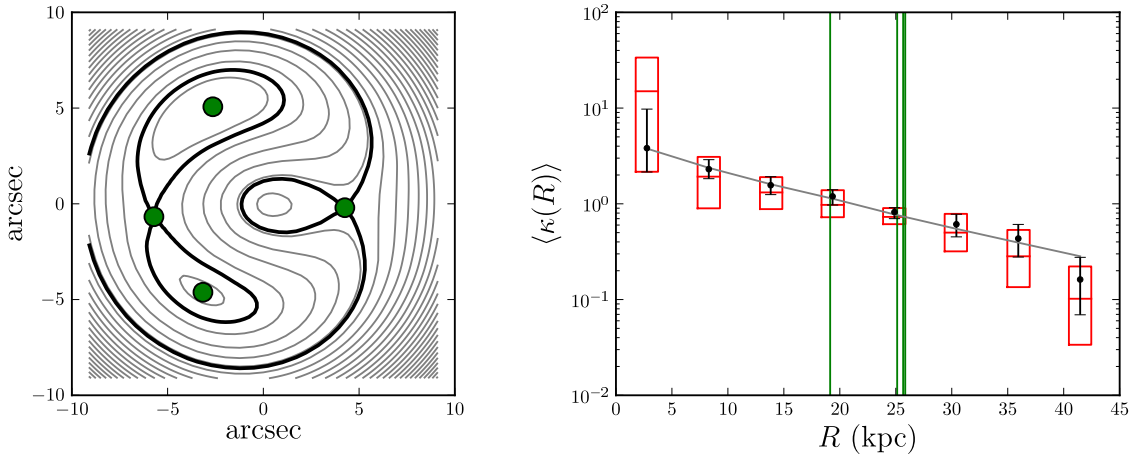

Figure 8. The reconstructed lens model of a source producing four images. **(Left)** The arrival time surface of the lens. The images **(dots)** are at the stationary points of the surface. **(Right)** The radial density profile of the original lens mass **(grey)** compared to the ensemble of possible models produced from the MCMC sampling **(black error bars)**. The boxes **(red)** show the results obtained from PixeLens for the same lens where the error estimates clearly favor shallower models suggesting that the old algorithm over-samples some regions of the parameter space. The radial image positions are marked by vertical lines.

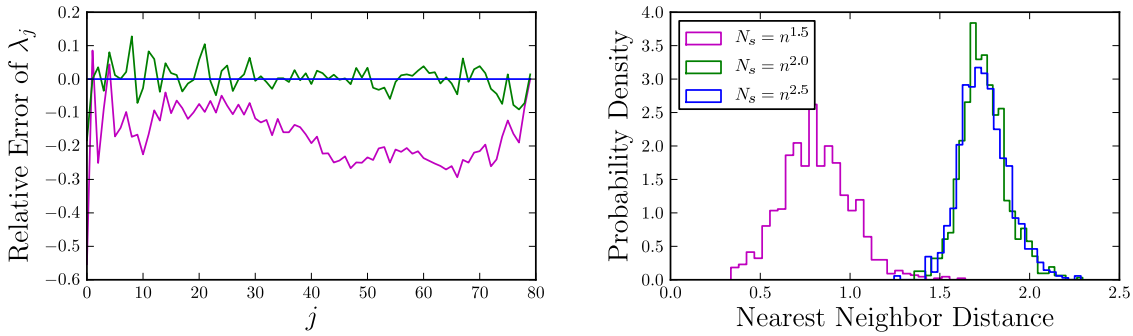

Figure 9. A convergence test of the sample eigenvalues **(left)** and the probability of nearest neighbor distances **(right)** using the realistic lens model depicted in Figure 8. The number of steps N_s used in the MCMC walk is increased until the histograms of nearest neighbor distances converge. An accept/reject random sample is impossible to obtain given the unknown shape of the simplex. The left plot has been normalized to the $N_s = n^{2.5}$ case. For $N_s \gtrsim n^2$ the eigenvalues are stable.

Table 1. Estimated volumes for the tested simplices of H_n, S_n, and the gravitational lens for various values of N_s, demonstrating convergence. The volumes are calculated from the eigenvalues shown in Figure 6 and Figure 9. The column labeled random is calculated via direct random sampling, which is not possible for the lens. Where the value is red, the volume is underestimated compared with the random sample.

	$n^{1.5}$	n^2	$n^{2.5}$	n^3	Random	
H_n	0.244 ± 0.044	9.1 ± 1.5	8.9 ± 1.4	—	9.0 ± 1.4	$\times 10^{-214}$
S_n	—	0.094 ± 0.043	1.64 ± 0.51	1.62 ± 0.46	1.63 ± 0.52	$\times 10^{-202}$
Lens	$(9.5 \pm 8.5) \times 10^{-8}$	0.94 ± 0.19	1.21 ± 0.24	—	—	$\times 10^{-93}$

ple models cannot. For comparison, we also show the results for the same lens obtained using PixeLens (red boxes). Although the results are similar, the PixeLens error estimates favor shallower or nearly flat models, again suggesting that the old algorithm over-samples some regions of the parameter space as shown in the upper-right panel of Figure 2.

We also performed the same eigenvalue and nearest neighbor analysis as before but because we are unable to directly sample the solution space we can only change N_s until we converge. As expected and shown in Figure 9 we converge when $N_s \sim n^2$.

7 OUTLOOK

Our novel proposal strategy for the Metropolis-Hastings algorithm allows sampling of general convex polytopes in 100 or more dimensions. We have implemented an efficient parallel version of the algorithm in the gravitational lens modeling framework GLASS so that we may apply the strategy to large lensing problems exceeding 1000 dimensions. GLASS will be publicly available for download in the near future.

Several future applications are possible. Multiple redshift sources carry information of the cosmological distances, which in turn depend on the cosmological parameters. Considering the statistical dispersion of the parameter space, one could in principle be able to infer the cosmological parameters in a Bayesian framework. In order to achieve this, a uniform sample of the solution space is needed.

Previous work on estimating the Hubble Time has used systems of up to eighteen lenses (Paraficz & Hjorth 2010). New lenses can be included to further constrain this value but each additional lens increases the dimensionality of the space by ~ 100 making this current work essential for such upcoming projects.

ACKNOWLEDGEMENTS

The authors would like to thank Prasenjit Saha for careful reading of the manuscript and extremely helpful advice.

REFERENCES

Auger M. W., Treu T., Bolton A. S., Gavazzi R., Koopmans L. V. E., Marshall P. J., Bundy K., Moustakas L. A., 2009, ApJ, 705, 1099
Avis D., Fukuda K., 1992, Discrete & Computational Geometry, 8, 295
Bartelmann M., Schneider P., 2001, Phys. Rep., 340, 291
Belisle C. J. P., Romeijn H. E., Smith R. L., 1993, Mathematics of Operations Research, 18, pp. 255
Blandford R., Narayan R., 1986, ApJ, 310, 568
Cappellari M. et al., 2006, MNRAS, 366, 1126
Chib S., Greenberg E., 1995, The American Statistician, 49, pp. 327
Clowe D., Bradač M., Gonzalez A. H., Markevitch M., Randall S. W., Jones C., Zaritsky D., 2006, ApJ, 648, L109
Coe D., Fuselier E., Benítez N., Broadhurst T., Frye B., Ford H., 2008, ApJ, 681, 814
Coles J., 2008, ApJ, 679, 17
Dantzig G. B., 1963, Linear Programming and Extensions. Princeton University Press, Princeton, NJ
Devroye L., 1986, Non-Uniform Random Variate Generation. Springer-Verlag
Dyer M. E., 1983, Mathematics of Operations Research, 8, pp. 381
Faure C. et al., 2011, A&A, 529, A72
Faure C. et al., 2008, ApJS, 176, 19
Gelman A., Roberts G. O., Gilks W. R., 1996, in Bayesian statistics, 5 (Alicante, 1994), Oxford Sci. Publ., Oxford Univ. Press, New York, pp. 599–607
Hastings W. K., 1970, Biometrika, 57, pp. 97
Kochanek C. S., Falco E. E., Impey C. D., Lehár J., McLeod B. A., Rix H.-W., 1999, in American Institute of Physics Conference Series, Vol. 470, After the Dark Ages: When Galaxies were Young (the Universe at $2 < Z < 5$), Holt S., Smith E., eds., pp. 163–175
Koopmans L. V. E., 2005, MNRAS, 363, 1136
McMullen P., Shephard G. C., 1971, Convex polytopes and the upper bound conjecture. Cambridge University Press, London
Merritt D., 1999, PASP, 111, 129
Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, J. Chem. Phys., 21, 1087
Oguri M., 2007, ApJ, 660, 1
Paraficz D., Hjorth J., 2010, ApJ, 712, 1378
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 2007, Numerical Recipes 3rd Edition: The Art of Scientific Computing, 3rd edn. Cambridge University Press, New York, NY, USA
Read J. I., Saha P., Macciò A. V., 2007, ApJ, 667, 645
Refsdal S., 1964a, MNRAS, 128, 307
Refsdal S., 1964b, MNRAS, 128, 295
Robert C. P., Casella G., 2005, Monte Carlo Statistical

Methods (Springer Texts in Statistics). Springer-Verlag New York, Inc., Secaucus, NJ, USA
Saha P., Coles J., Macciò A. V., Williams L. L. R., 2006, ApJ, 650, L17
Saha P., Williams L. L. R., 2004, AJ, 127, 2604
Schneider P., 2006, in Saas-Fee Advanced Course 33: Gravitational Lensing: Strong, Weak and Micro, Meylan G., Jetzer P., North P., Schneider P., Kochanek C. S., Wambsganss J., eds., pp. 1–89
Schneider P., Ehlers J., Falco E. E., 1992, Gravitational Lenses, Schneider, P., Ehlers, J., & Falco, E. E., ed.
Schwarzschild M., 1979, ApJ, 232, 236
Schwarzschild M., 1982, ApJ, 263, 599
Sereno M., Lubini M., Jetzer P., 2010, A&A, 518, A55
Tortora C., Napolitano N. R., Romanowsky A. J., Jetzer P., 2010, ApJ, 721, L1
Young P., Gunn J. E., Oke J. B., Westphal J. A., Kristian J., 1981, ApJ, 244, 736

This paper has been typeset from a TEX/ LATEX file prepared by the author.

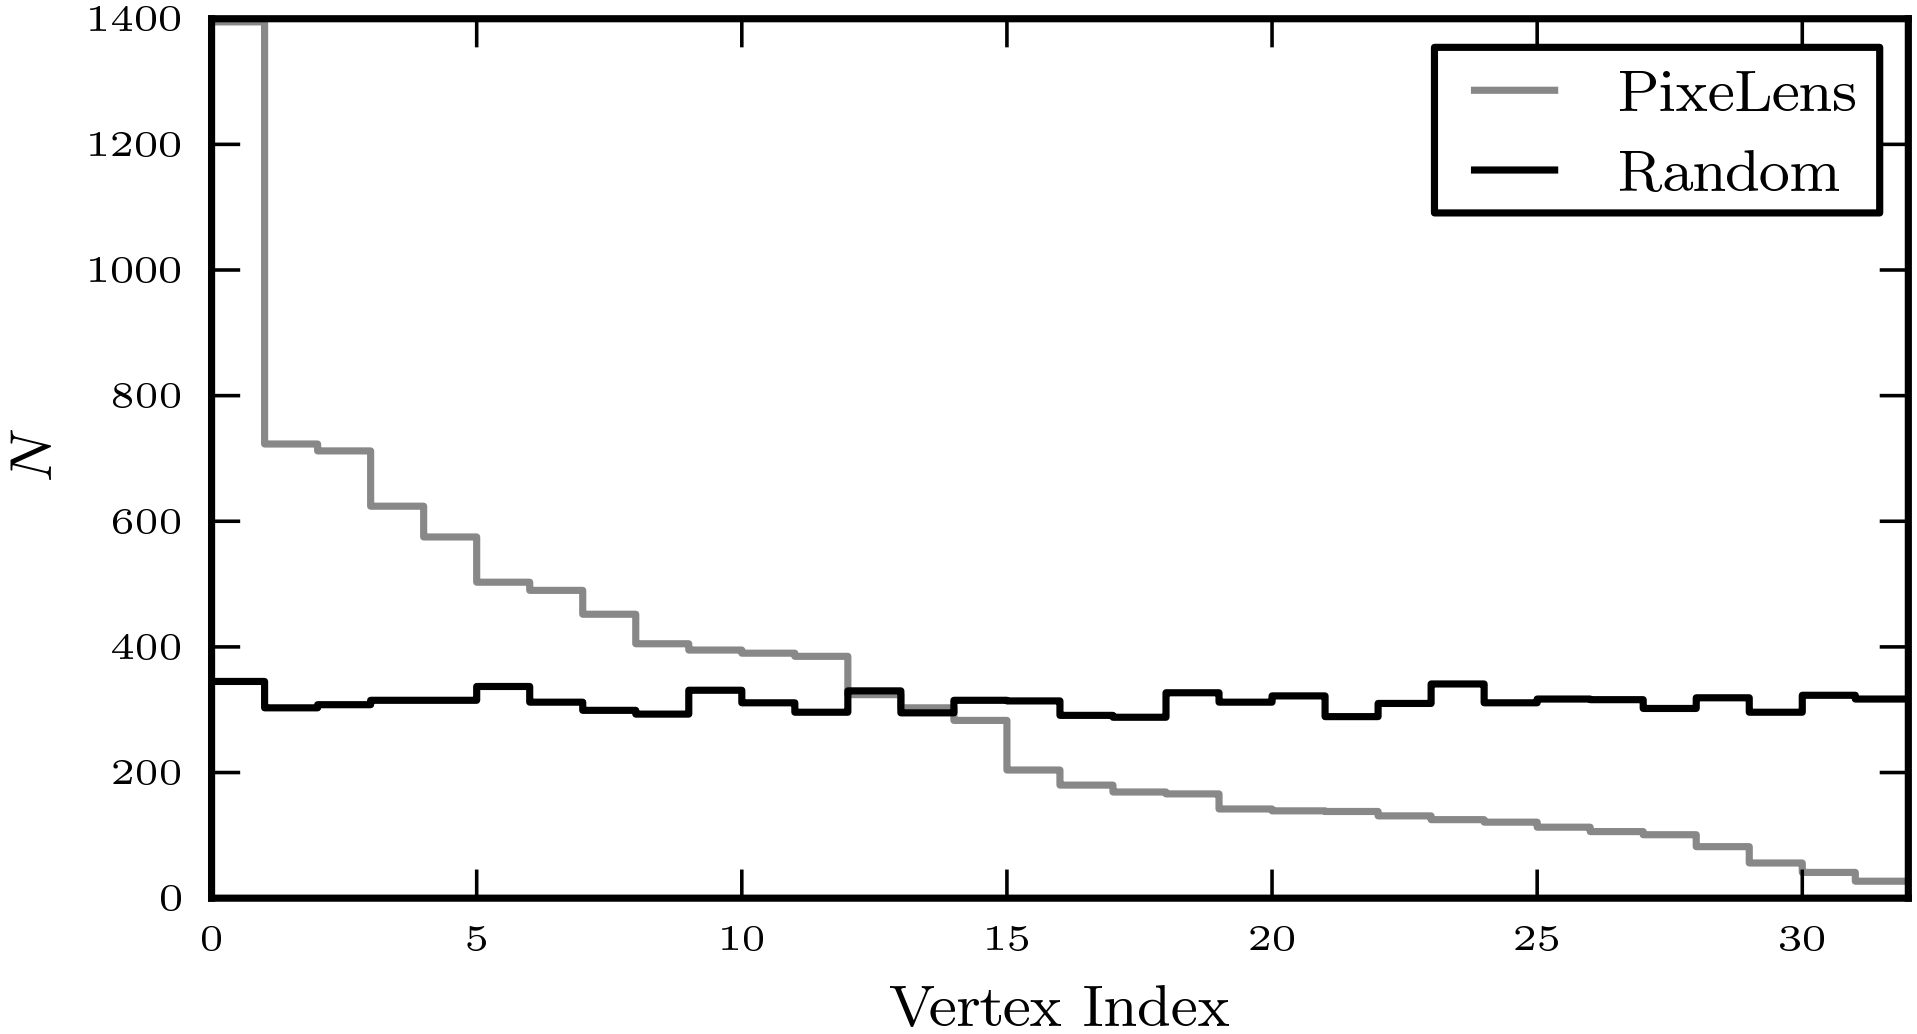
1400
1200
1000
800
600
400
200
0
N
0
5
10
15
20
25
30
Vertex Index
PixeLens
Random

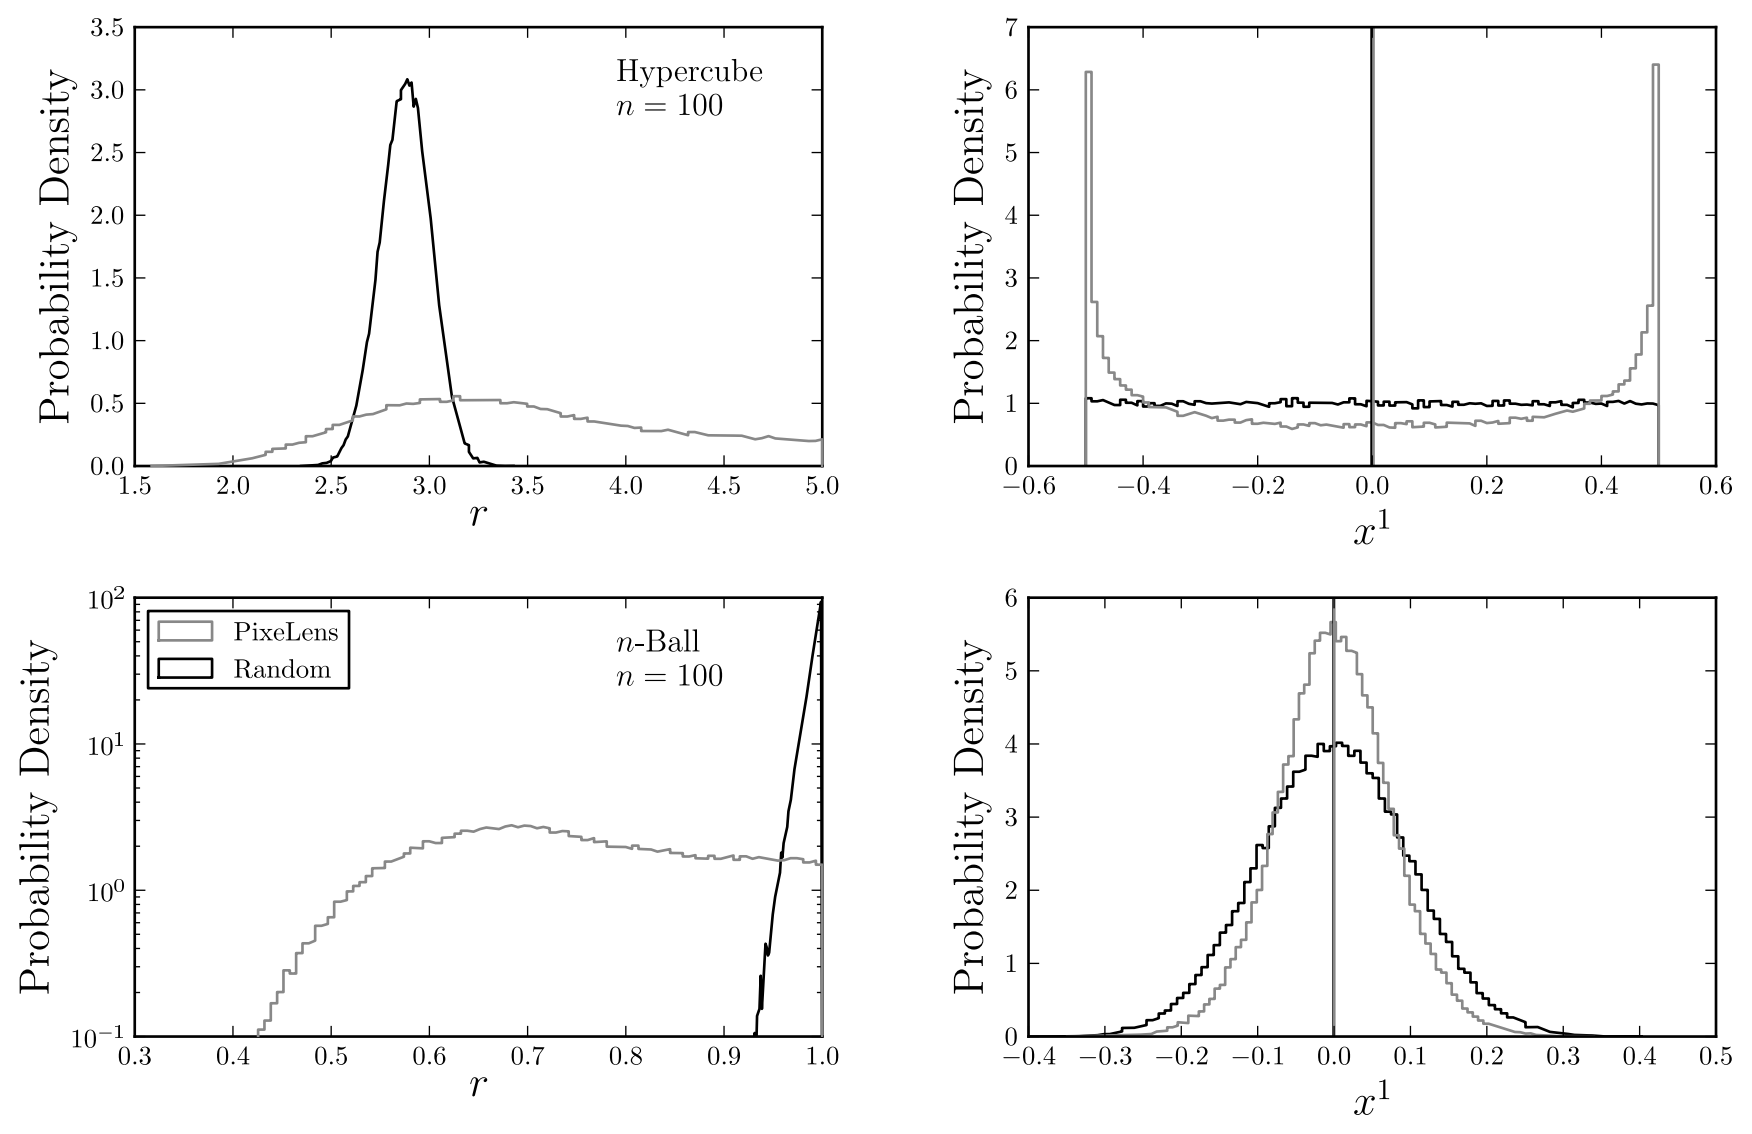
Hypercube
$n = 100$
Probability Density
r
Probability Density
x^1
PixeLens
Random
n-Ball
$n = 100$
Probability Density
r
Probability Density
x^1

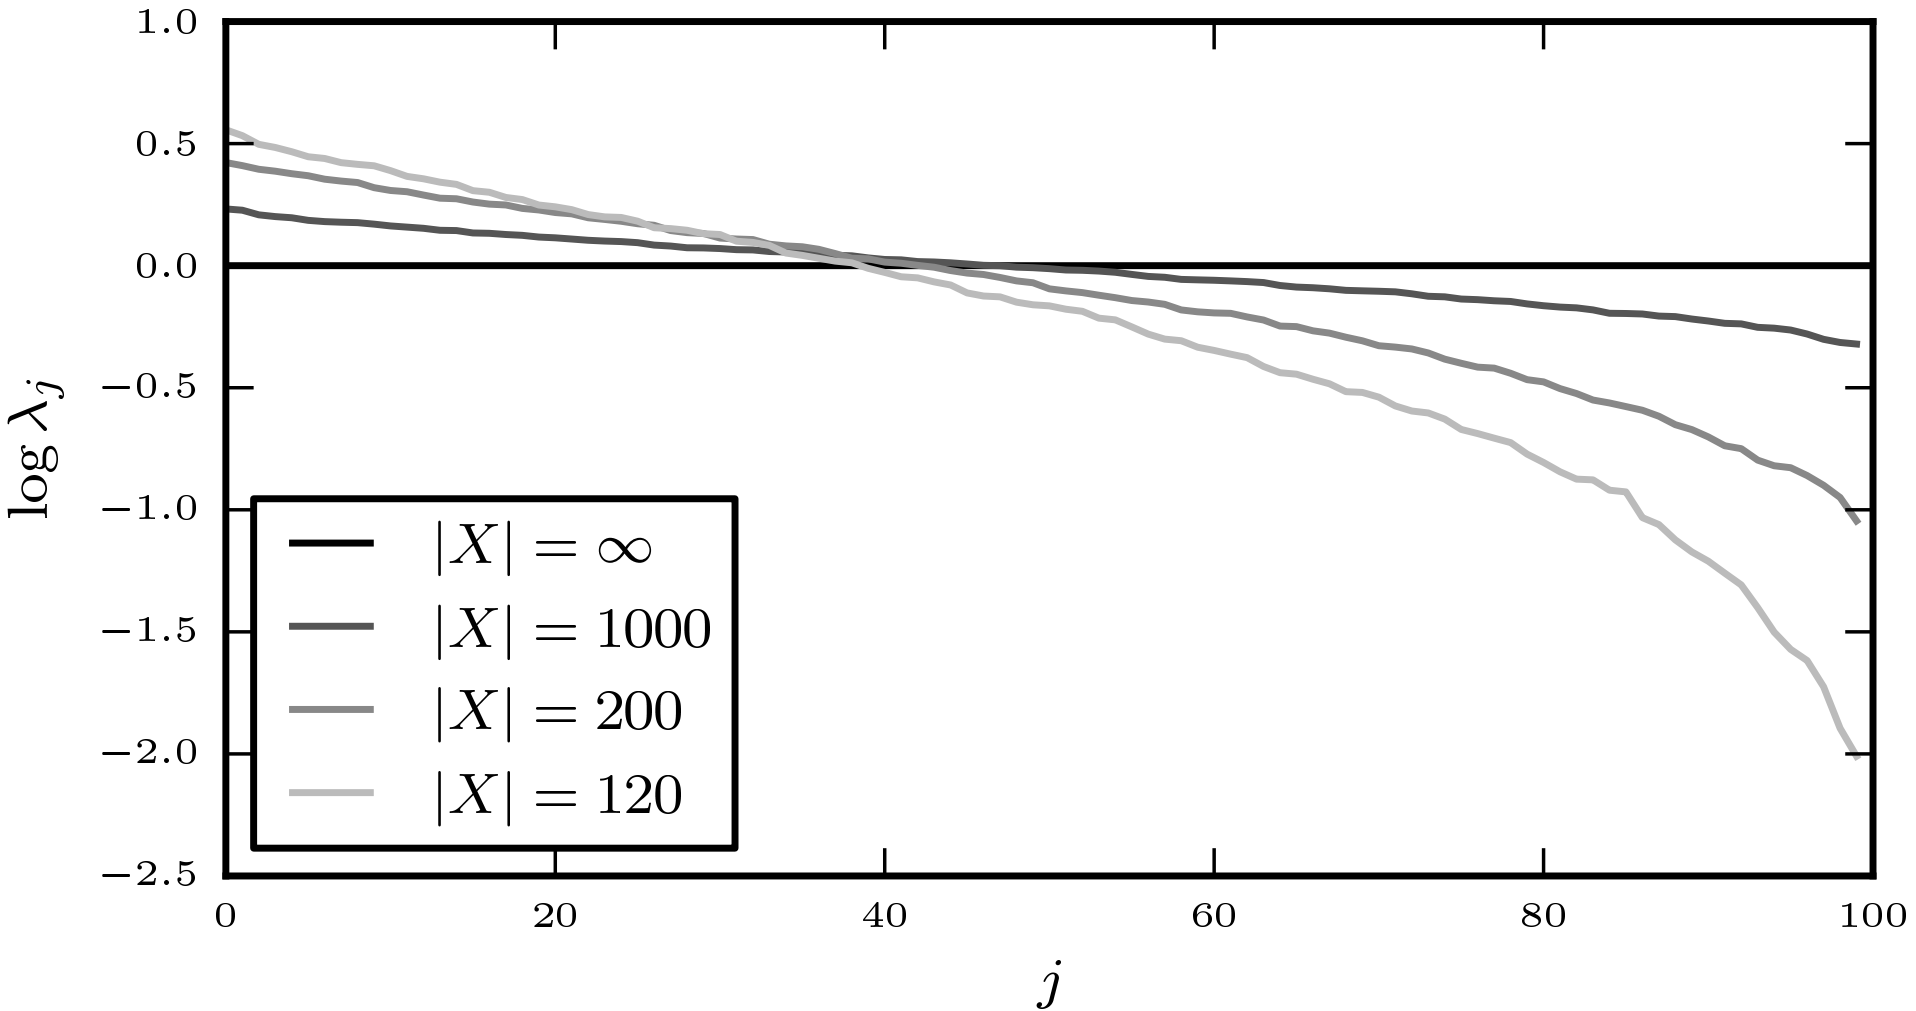
1.0
0.5
0.0
−0.5
−1.0
−1.5
−2.0
−2.5
log λ_j
0
20
40
60
80
100
j
|X| = ∞
|X| = 1000
|X| = 200
|X| = 120

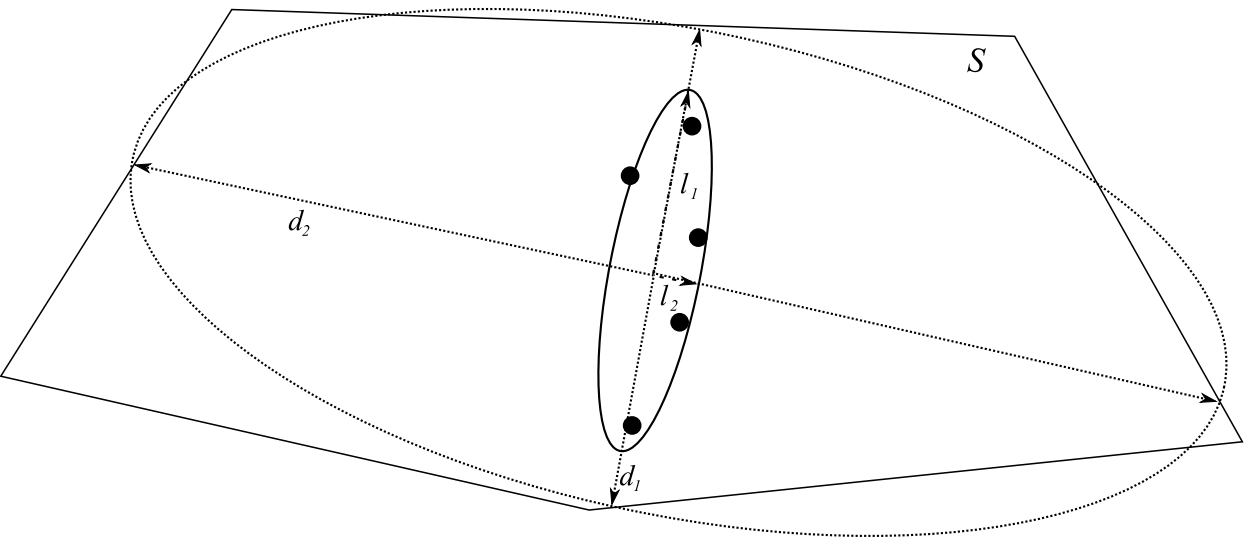
S
l_1
l_2
d_2
d_1

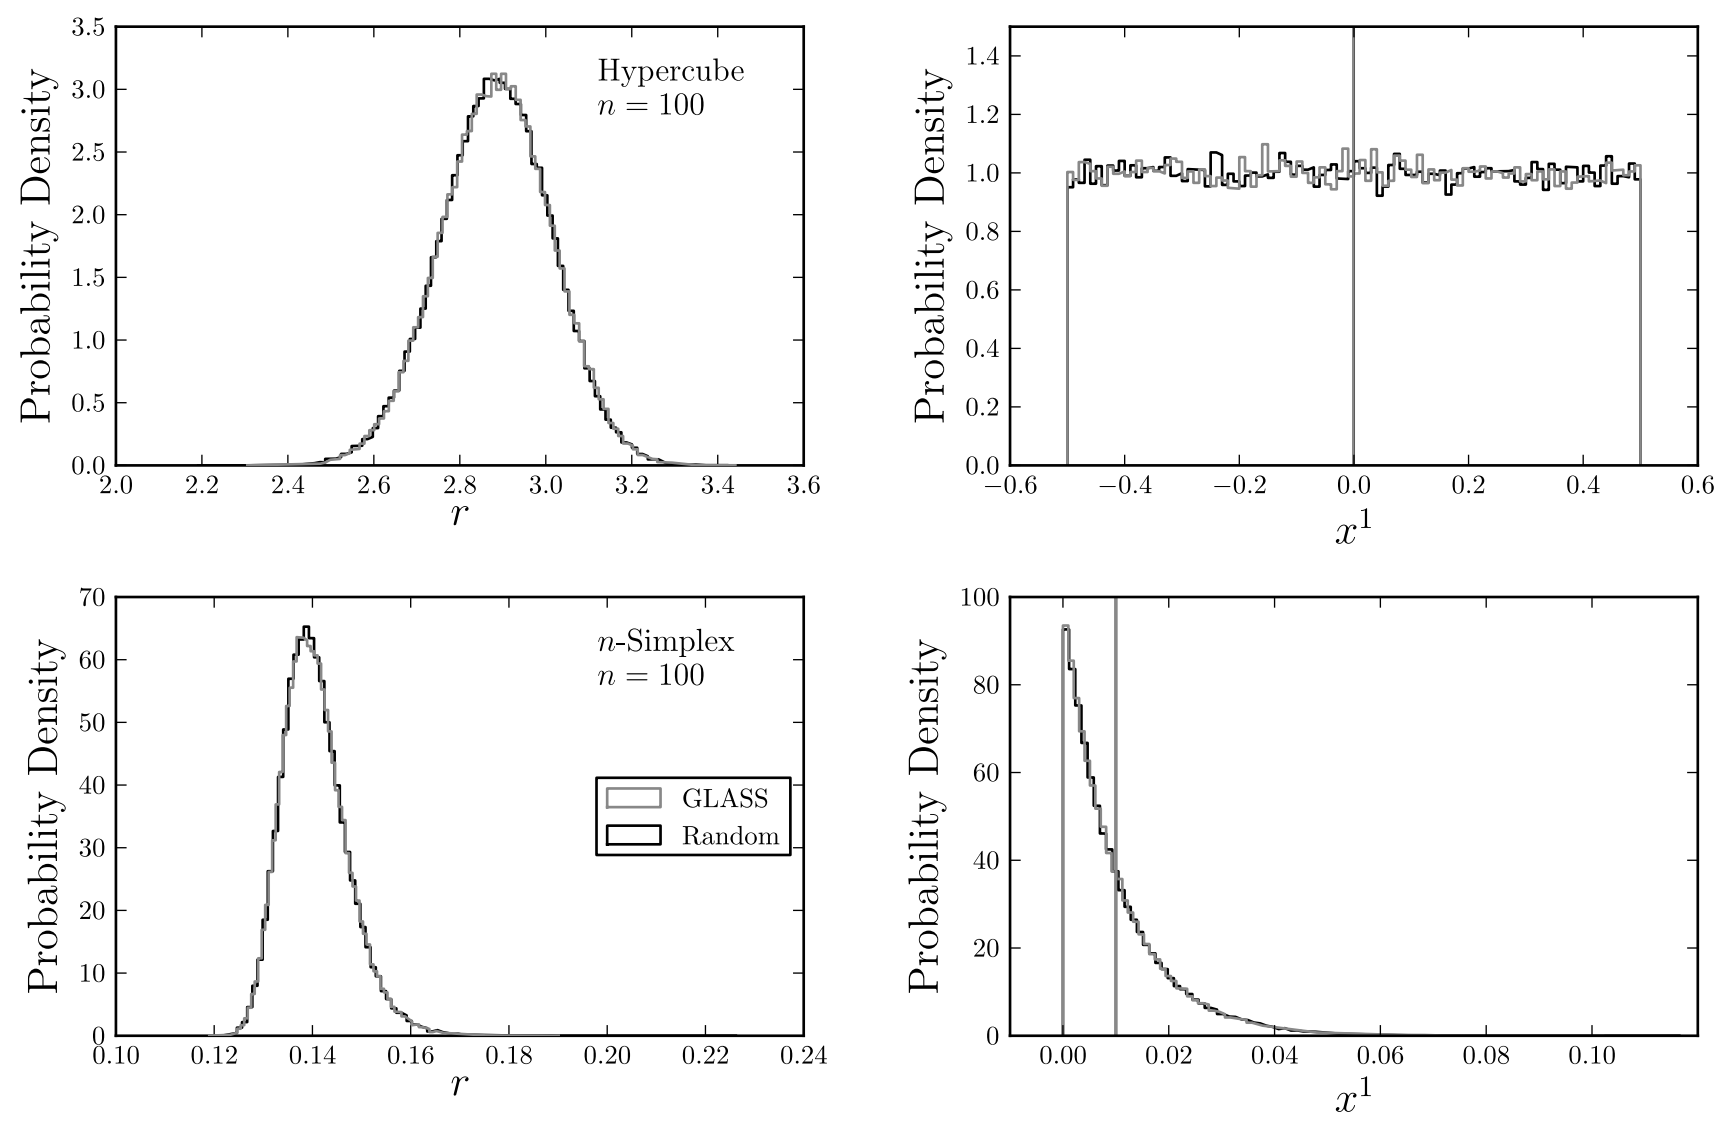
Hypercube
$n = 100$
Probability Density
r
Probability Density
x^1
n-Simplex
$n = 100$
GLASS
Random
Probability Density
r
Probability Density
x^1

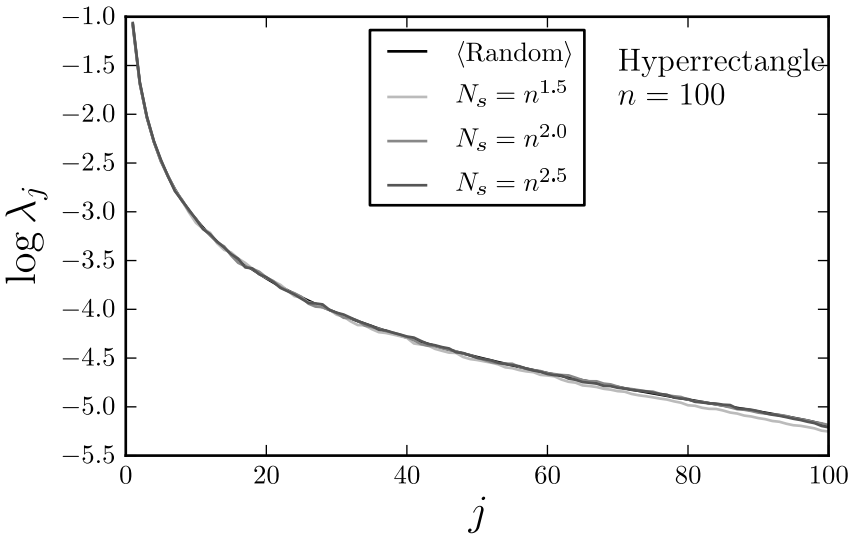
⟨Random⟩
$N_s = n^{1.5}$
$N_s = n^{2.0}$
$N_s = n^{2.5}$
Hyperrectangle
$n = 100$
$\log \lambda_j$
j

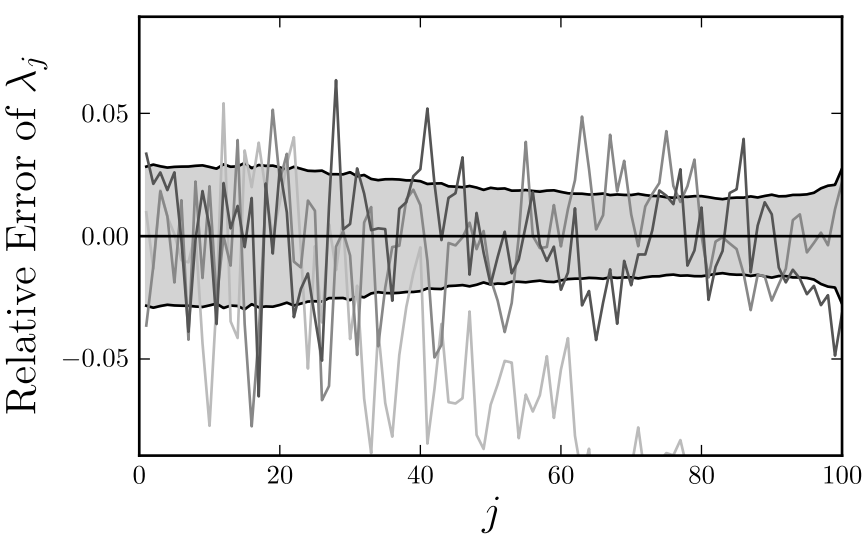
Relative Error of λ_j
j

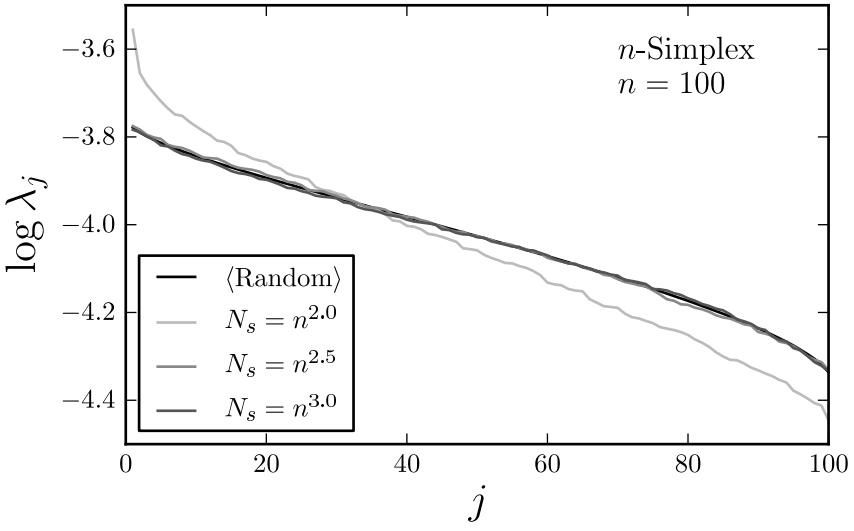
⟨Random⟩
$N_s = n^{2.0}$
$N_s = n^{2.5}$
$N_s = n^{3.0}$
n-Simplex
$n = 100$
$\log \lambda_j$
j

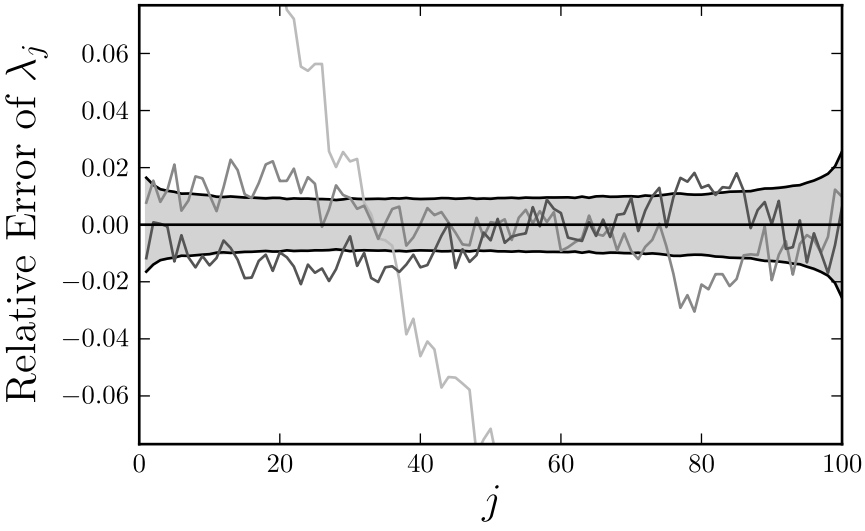
Relative Error of λ_j
j

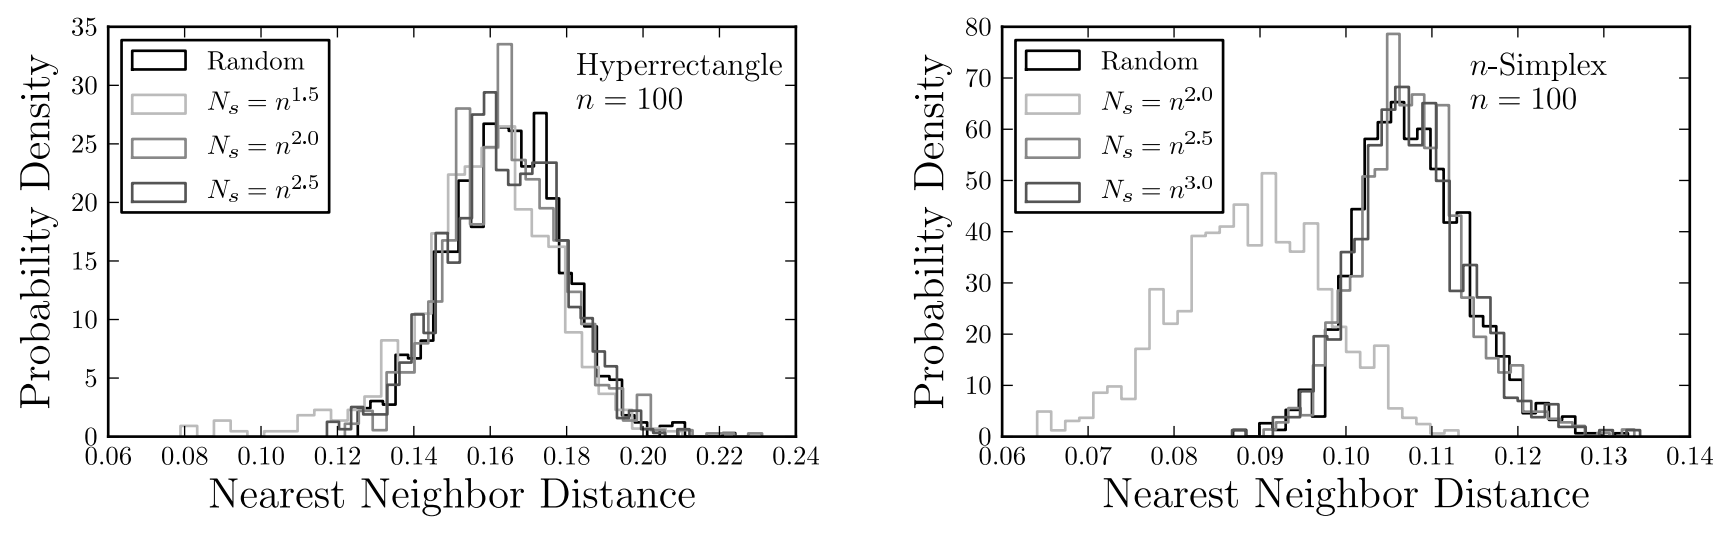
Random
$N_s = n^{1.5}$
$N_s = n^{2.0}$
$N_s = n^{2.5}$
Hyperrectangle
$n = 100$
Probability Density
Nearest Neighbor Distance
0.06
0.08
0.10
0.12
0.14
0.16
0.18
0.20
0.22
0.24
0
5
10
15
20
25
30
35
Random
$N_s = n^{2.0}$
$N_s = n^{2.5}$
$N_s = n^{3.0}$
n-Simplex
$n = 100$
Probability Density
Nearest Neighbor Distance
0.06
0.07
0.08
0.09
0.10
0.11
0.12
0.13
0.14
0
10
20
30
40
50
60
70
80

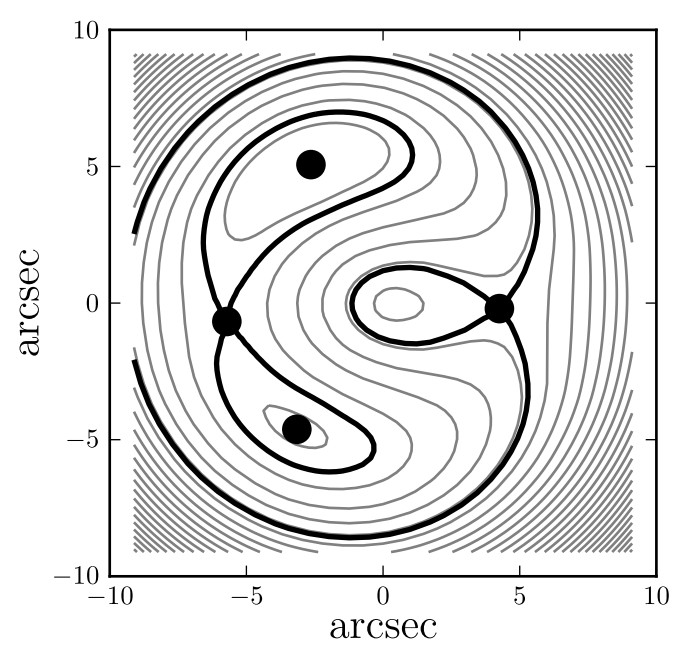
arcsec
arcsec

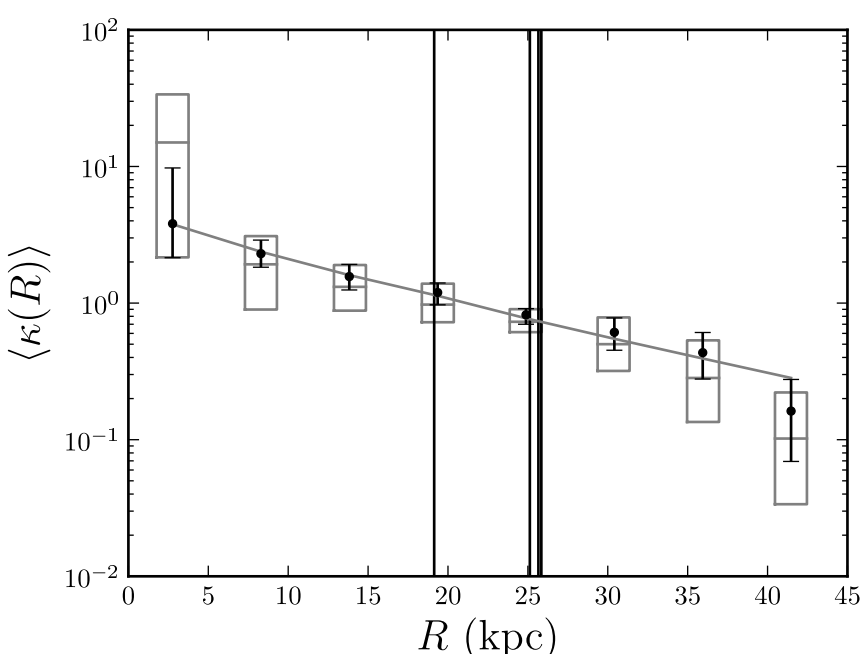
$\langle \kappa(R) \rangle$
R (kpc)

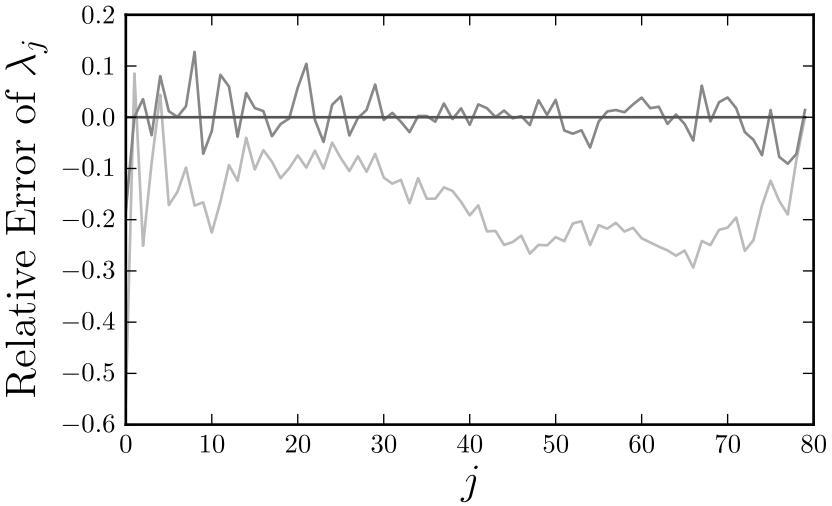
Relative Error of λ_j
j
0.2
0.1
0.0
−0.1
−0.2
−0.3
−0.4
−0.5
−0.6
0
10
20
30
40
50
60
70
80

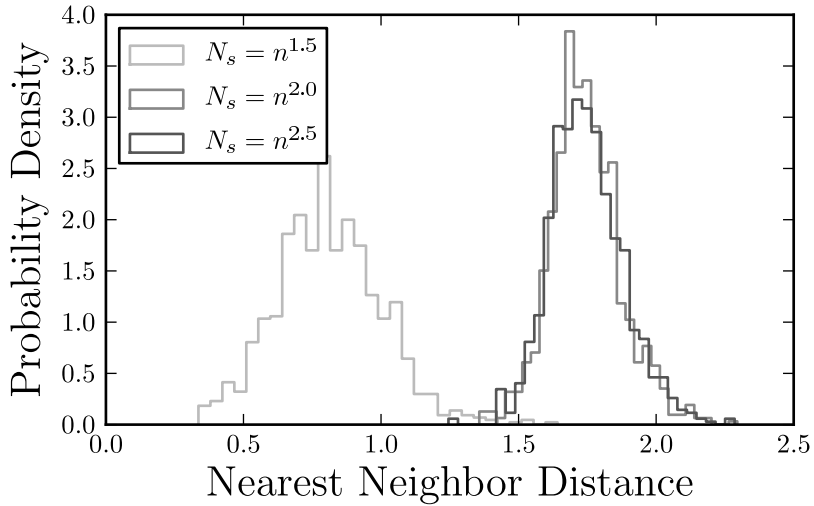
Probability Density
Nearest Neighbor Distance
$N_s = n^{1.5}$
$N_s = n^{2.0}$
$N_s = n^{2.5}$
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
0.0
0.5
1.0
1.5
2.0
2.5